

Efficient Desynchronization of Thermostatically Controlled Loads

Jan Bendtsen ^{*,*} Srinivas Sridharan ^{**}

^{*} *Department of Electronic Systems, Automation and Control, Aalborg University, Denmark (e-mail: dimon@es.aau.dk).*

^{**} *Department of Mechanical and Aerospace Engineering, University of California, San Diego, California (email: srsridharan@eng.ucsd.edu)*

Abstract: This paper considers demand side management in smart power grid systems containing significant numbers of thermostatically controlled loads such as air conditioning systems, heat pumps, etc. Recent studies have shown that the overall power consumption of such systems can be regulated up and down centrally by broadcasting small setpoint change commands without significantly impacting consumer comfort. However, sudden simultaneous setpoint changes induce undesirable power consumption oscillations due to sudden synchronization of the on/off cycles of the individual units. In this paper, we present a novel algorithm for counter-acting these unwanted oscillations, which requires neither central management of the individual units nor communication between units. We present a formal proof of convergence of homogeneous populations to desynchronized status, as well as simulations that indicate that the algorithm is able to effectively dampen power consumption oscillations for both homogeneous and heterogeneous populations of thermostatically controlled loads.

1. INTRODUCTION

With growing penetration of renewable energy sources in modern power grids, *demand side management* has been gaining attention as a means of achieving better balancing between supply and demand (Mohsenian-Rad et al. [2010], Strbac [2008], Short et al. [2007]). Indeed, it appears that the higher intermittency and lack of dispatchability associated with increased dependence on renewable energy sources can be taken care of more effectively by electrical loads than by conventional generators, which are typically not designed for fast up- and down regulation (Strbac [2008], Klobasa [2010]).

Various technologies are currently being considered in the context of demand side management; coordinated charging of batteries, e.g., in electric vehicles (Mets et al. [2010]), deliberate scheduling of loads with flexible deadlines (Petersen et al. [2012]) as well as allowing local consumers with slow dynamics (large time constants) to store more or less energy at convenient times and thereby adjusting the momentary consumption (e.g., Moslehi and Kumar [2010]), among others. In particular, so-called *thermostatically controlled loads* (TCLs), such as deep freezers, refrigerators, local heat pumps etc., are showing great potential in this context, since they account for a large volume of consumption in most countries with significant renewable penetration—for instance, as of 2009, about 87 percent of all US homes were equipped with air conditioning.¹

Thus, at least in theory, manipulating the operating conditions of large populations of units slightly while avoiding discomfort to end users appears to be attractive, given that it can be achieved simply by broadcasting setpoint changes. Control strategies based on this principle were considered in Callaway [2009], Kundu et al. [2011], Perfumo et al. [2012] and Bashash and Fathy [2012], among others. In these strategies, subtle changes in the thermostat setpoint temperature (less than a degree Celcius) are transmitted to all the participating consumers simultaneously. Such small variations in the thermostat setpoints are expected to remain almost unnoticed by the customers, but if a sufficiently large number of units are shifted in the same direction, the overall power consumption can be shifted quickly by significant amounts, allowing to compensate for power production fluctuations in the grid.

There is an intrinsic problem with this approach, however. When the setpoint is changed concurrently in a large number of TCLs, their on/off cycles tend to become *synchronized*, which leads to large unwanted fluctuations (damped oscillations) in the overall power consumption Callaway [2009], Kundu et al. [2011]. If the entire population of TCLs consists of identical units (*homogeneous*), it is possible to counteract the fluctuations by means of a centralized control law. In the much more realistic case of non-identical (*heterogeneous*) units, on the other hand, the fluctuations are far harder to remove using centralized control strategies. To the authors' knowledge, the only solution strategy presented in the literature so far is Kundu and Sinitsyn [2012], where a strategy in which the individual TCLs deliberately slow down their transition in order to avoid synchronization was proposed.

In this paper, we propose a novel, decentralized algorithm for avoiding synchronization without having to ac-

^{*} The first author was supported in part by the Otto Mønsted and Dir. Ib Henriksen foundations, Denmark; The second author was partially supported by AFOSR.

¹ <http://www.eia.gov/consumption/residential/reports/2009/air-conditioning.cfm>

tively communicate between units. The algorithm proposed here is fundamentally different from the one presented in Kundu and Sinitsyn [2012]) in that it first adjusts the temperature bands in which the TCLs operate to achieve early (randomized) desynchronization, and then makes use of a technique inspired by contention-based media access protocols to adaptively adjust the on/off cycles to achieve as little synchronization among units as possible. We present a formal proof of convergence of homogeneous populations to this desynchronized status, as well as simulations that indicate that the algorithm is able to effectively dampen power consumption oscillations for both homogeneous and heterogeneous populations of thermostatically controlled loads.

The outline of the rest of the paper is as follows. Section 2 first presents the simplified TCL model we shall employ, along with the synchronization issue. Section 3 then presents the proposed algorithm along with the aforementioned convergence result. Section 4 illustrates the feasibility of the approach through two simulation examples. Finally, Section 5 offers some concluding remarks.

2. PROBLEM SETUP

We consider a population of N individual temperature controlled loads. The TCLs are modeled in a generic manner, which is deliberately kept as simple as possible. Nonetheless, as will be illustrated with a few simulation examples, the behavior of a large population of simple units can be fairly complex.

2.1 Model

The individual TCLs are modeled as follows. Let the internal and ambient temperatures of the volumes (living spaces, refrigerators, cold storages etc.) affected by the action of the heating/cooling hardware of the i 'th consumer be denoted θ_i and $\theta_{\infty,i}$, respectively, and assume that the hardware is purely on/off-regulated. A simple model for the internal temperature can then be formulated as Malhame and Chong [1985], Callaway [2009], Bashash and Fathy [2012]:

$$\dot{\theta}_i(t) = \frac{1}{R_i C_i} (\theta_{\infty,i} - \theta_i(t) - s_i(t) R_i P_i), \quad (1)$$

$$s_i(t) = \begin{cases} 0 & \text{if } s_i(t^-) = 1 \wedge T_i(t) \leq T_{\min,i} \\ 1 & \text{if } s_i(t^-) = 0 \wedge T_i(t) \geq T_{\max,i} \\ s_i(t^-) & \text{otherwise} \end{cases} \quad (2)$$

for $i = 1, 2, \dots, N$, where C_i is the thermal capacitance of the consumer, R_i is the corresponding thermal resistance and P_i is the (constant) heating/cooling power supplied by the hardware when switched on. $s_i \in \{0, 1\}$ is a binary switching variable that determines whether or not the hardware is turned on; basically, it switches status whenever the internal temperature encounters the limits of a pre-set temperature span $[\theta_{\min,i}, \theta_{\max,i}] \subset \mathbb{R}$ ($t^- = t - \epsilon$, where ϵ is an infinitesimal positive value). Note that both heating and cooling systems can be modeled by (4)–(2), by choosing the sign of the term $s_i(t) R_i P_i$ appropriately.

The temperature limits $\theta_{\min,i}$ and $\theta_{\max,i}$ are related to the i 'th consumer's setpoint $\theta_{\text{sp},i}$ through the relations




Fig. 1. Temperature and power consumption of a single TCL

$$\theta_{\min,i} = \theta_{\text{sp},i} - \frac{\Delta}{2}, \quad \theta_{\max,i} = \theta_{\text{sp},i} + \frac{\Delta}{2}$$

where Δ is the width of the temperature interval. Furthermore, the cumulative power consumption of the population of TCLs at any given time t can be computed as

$$P(t) = \sum_{i=1}^N \frac{P_i s_i(t)}{\eta_i} \quad (3)$$

where η_i is the coefficient of performance for the i 'th heating/cooling unit.

Figure 1 illustrates the characteristic behavior of a single TCL. The parameters are chosen as $R = 2^\circ\text{C/kW}$, $C = 5 \text{ kWh}/^\circ\text{C}$, $P = 14 \text{ kW}$, $\eta = 2.5$, $\theta_{\text{sp}} = 20^\circ\text{C}$, $\Delta = 1^\circ\text{C}$ and $\theta_{\infty} = 28^\circ\text{C}$, respectively. The temperature setpoint is changed to $\theta_{\text{sp}} = 20.5^\circ\text{C}$ at time $t = 10 \text{ h}$, resulting in a slightly longer period in which the cooling power is turned off, until the new operating interval is reached.

2.2 Population behavior

As illustrated in Callaway [2009], Perfumo et al. [2012], Bashash and Fathy [2012] and others, the power consumption of populations of units with the individual dynamics given above can be manipulated via broadcasts of small setpoint changes. Figure 2 shows the behavior of a population of 10,000 homogeneous (identical) TCLs with the same set of parameters as in figure 1 (only the temperature profiles of 25 TCLs are shown in the figure). The devices are started at random temperatures uniformly distributed within $[\theta_{\min}, \theta_{\max}]$, and their on/off cycles are thus desynchronized, yielding a fairly smooth power consumption trajectory. At time $t = 10 \text{ h}$, all the TCLs are subjected to a common step in the setpoint of 0.5°C . As can be seen, the overall power consumption drops by an amount proportional to the size of the step, but significant oscillations appear in the power consumption due to synchronization of the individual units' on/off cycles.

Figure 3 illustrates the behavior of a population of TCLs under the same simulation circumstances, but with C_i chosen from a normal distribution with mean 5 and spread $0.5 \text{ kWh}/^\circ\text{C}$ (again, only the temperature profiles of 25 TCLs are shown in the figure). A step of 0.5°C in the setpoint is broadcast to all the TCLs at time $t = 10 \text{ h}$. Notice how the power consumption is quite smooth before the step




Fig. 2. Temperature and power consumption of a population of 10,000 homogeneous TCLs




Fig. 3. Temperature and power consumption of a population of 10,000 heterogeneous TCLs

because the on/off cycles of the TCLs are desynchronized, whereas immediately after the step the TCLs become synchronized and the power consumption again exhibits large fluctuations, which in this case die out slowly. The oscillations die out faster if the parameters vary more (e.g., if C_i is chosen from a distribution with larger spread); furthermore, the amplitude is relatively larger for larger populations of units. However, qualitatively the behavior remains the same. The oscillations are known as *parasitic oscillations* and are hard to remove via centralized control signals (Kundu and Sinitsyn [2012]).

2.3 Problem formulation

In order to reduce the aforementioned oscillations in power consumption after the setpoint change broadcasts, it is clearly necessary to deliberately desynchronize the TCLs without interrupting their operation. However, a centralized algorithm for doing so is undesirable; it is considered infeasible to keep close track of the internal states (temperatures, set points etc.) of all the units in a centralized manner, since doing so would require regular measurement feedback from all the devices, which would give rise to a very heavy communication and computational load on the system.

Thus, we look for a low-complexity, *decentralized* algorithm that can be implemented locally in each TCL, and which satisfies the following requirements:

- R1 The algorithm may only use local information; this means that the TCLs may not communicate with each other, and labels identifying each individual TCL may not be pre-assigned
- R2 Communication with the power supply utility must be limited to broadcast from the utility; no communication originating from the TCLs is allowed
- R3 The general operation of each individual TCL may not be altered in a manner that is detrimental to the user comfort; for example, the unit may not be deliberately kept turned off or on long enough for the temperature to leave the interval $[\theta_{\min}, \theta_{\max}]$

3. DESYNCHRONIZATION ALGORITHM

In the following, we present a completely decentralized algorithm that is able to reduce the power oscillations discussed above without violating any of the requirements R1–R3. The only significant assumption we require to be satisfied is that the individual TCLs have access to time stamps of changes in the total power consumption. While this may be considered a form of global information, its usage is limited to identifying the times at which the immediate predecessor and the immediate successor of a local TCL changes state. It is not used to actually identify said units; i.e., we only require local temporal information. The information about changes in global power consumption may either be measured locally or broadcast from the utility, thus avoiding violating requirements R1 and R2.

The algorithm comprises two main components. Firstly, the algorithm is reset whenever each TCL in question receives a setpoint change $\theta_{\text{sp},i} \rightarrow \theta_{\text{sp},i} + \delta$. At this point, each TCL immediately narrows its operation interval by a random value $\alpha \in [0, \Delta/2]$, causing the dynamics of unit i momentarily to be governed by the dynamics

$$\begin{aligned} \dot{\theta}_i(t) &= \frac{1}{R_i C_i} (\theta_{\infty,i} - \theta_i(t) - s_i(t) R_i P_i), \\ s_i(t) &= \begin{cases} 0 & \text{if } s_i(t^-) = 1 \wedge \theta_i(t) \leq \theta_{\text{sp},i} + \delta + \frac{\Delta_i}{2} - \alpha_i \\ 1 & \text{if } s_i(t^-) = 0 \wedge \theta_i(t) \geq \theta_{\text{sp},i} + \delta - \frac{\Delta_i}{2} + \alpha_i \\ s_i(t^-) & \text{otherwise} \end{cases} \end{aligned}$$

This operation rapidly induces a large degree of desynchronization among the population of TCLs without violating requirement R3, as illustrated in Figure 4. However, randomly choosing the temperature bounds may yield very tight bounds, which in turn may give rise to constant rapid on/off-switching. This is generally not desirable, so we drive the bounds back to their original settings by introducing the dynamics

$$\dot{\alpha}(t) = -a\alpha(t) \quad (4)$$

for some appropriately chosen $a > 0$, as also illustrated in the figure.

Now, even though most of the synchronization among the population of TCLs is likely to have disappeared with the random shrinking of the operating interval, it may still happen that some TCLs remain synchronized,




Fig. 4. Illustration of TCL desynchronization algorithm before and after a step occurring at $t = t_\delta$. The temperature profiles θ_i of three TCLs indexed by $i = 1, 2, 3$ are shown with full lines. The dashed lines indicate the operating intervals $[\theta_{sp,i} + \delta + \Delta_i/2 - \alpha_i(t), \theta_{sp,i} + \delta - \Delta_i/2 + \alpha_i(t)]$. Due to different random values of α_i , the TCLs become desynchronized after the step even though they were completely synchronized before and immediately after t_δ . The operation intervals converge exponentially back to their original values.

or close to synchronized. The second component of the proposed algorithm is designed to deal with any ‘left-over’ synchronization among TCLs that are close to being identical. Here, we make use of the fact that the behavior of the TCLs is periodic in steady state operation. Let the period be denoted T . Essentially, this part of the desynchronization algorithm repeatedly forces the devices to switch status at some point within the interval $[0, T]$, and adaptively adjusts this enforced switch timing to achieve maximal spread across the period; see Figure 5.

In short, each TCL records the switch timings of the immediately preceding and the immediately succeeding TCLs; then it sets the new enforced switch time to the mean of the switch times of the succeeding and preceding TCLs, i.e.,

$$t_{k+1} = \frac{t_{\text{prev}} + t_{\text{next}}}{2} \quad (5)$$

where $t_{k+1} \in [0, T]$ denotes the time of the enforced switch within the following period, modulo T .

The enforced switch timings are initially distributed randomly; the TCL which randomly obtains the smallest value will then not be able detect a switch prior to its own. When it recognizes this fact, it sets its own switch timing to 0 and maintains it at that value (modulo T); it is thus not necessary to assign a global label to the first and/or last TCL a priori.

Let x_k^j denote the time of the j ’th forced transition, modulo T , at iteration k ; that is, $0 = x_k^1 < x_k^2 < \dots < x_k^N < T$. We gather these switch timings in a vector $\mathbf{x}_k \in [0, T]^N$. We can then prove the following asymptotic result:

Theorem 1. If all switch timings for $i = 1, \dots, N$ are updated according to (5), the globally attracting switch

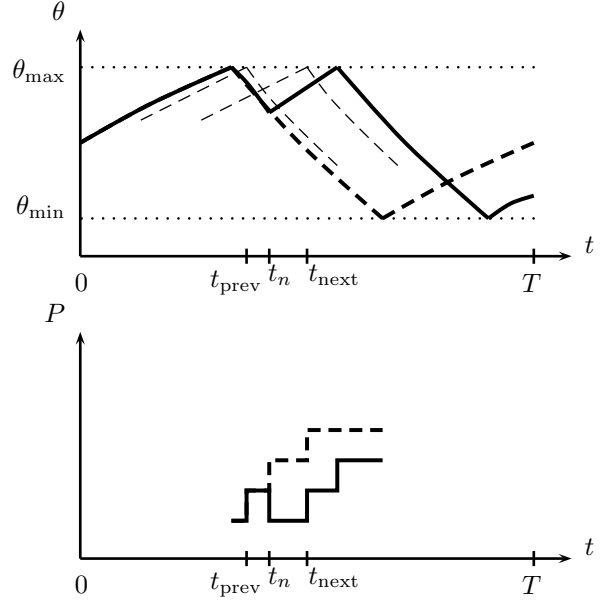


Fig. 5. Illustration of TCL desynchronization algorithm. The top figure shows the temperature profile of a TCL with (thick full line) and without (thick dashed line) enforced switching at time $t = t_n$. The thin dashed lines show parts of the temperature traces of the preceding and succeeding TCLs. Their switch times are recorded by the local TCL as t_{prev} and t_{next} , respectively. The bottom figure shows the relevant part of the corresponding power traces.

timing state \mathbf{x}^* (modulo T) is described by

$$\mathbf{x}^* \bmod T = \left[0, \frac{T}{N}, \frac{2T}{N}, \dots, \frac{T(N-1)}{N} \right]^T. \quad (6)$$

Proof: The algorithm described above for the evolution of the state (time) dynamics requires that each state be the mean value of the prior and next state with the caveat that the N ’th element should update its value based on the next pulse i.e. the value of the first state element in the next time instant. Hence, we first extend the state vector by placing the new value of the first state into the fictitious $N + 1$ th dimension of the state vector \mathbf{x} to generate a new state vector, denoted $\tilde{\mathbf{x}}$. The update dynamics of the enforced switch times can be written as

$$\begin{aligned} \tilde{\mathbf{x}}_{k+1} &= \begin{bmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 & \dots & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & \dots & 0 & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 \end{bmatrix} \tilde{\mathbf{x}}_k \\ &= \Gamma \tilde{\mathbf{x}}_k. \end{aligned} \quad (7)$$

which implies that

$$\lim_{k \rightarrow \infty} \tilde{\mathbf{x}}_k = \lim_{k \rightarrow \infty} \Gamma^k \tilde{\mathbf{x}}_0$$

Using eigenvalue decomposition, we may write the aforementioned limit as

$$\lim_{k \rightarrow \infty} \tilde{\mathbf{x}}_k = \lim_{k \rightarrow \infty} E \Lambda^k E^{-1} \tilde{\mathbf{x}}_0 \quad (8)$$

where $\Lambda \in \mathbb{R}^{(N+1) \times (N+1)}$ is a diagonal matrix of eigenvalues of Γ and $E \in \mathbb{R}^{(N+1) \times (N+1)}$ contains the corresponding eigenvectors. Γ has two repeated unit eigenvalues with linearly independent eigenvectors γ and $\mathbf{1} - \gamma$, where $\mathbf{1} \in \mathbb{R}^{N+1}$ is a vector of one-elements and

$$\gamma = \left[1 \ 1 - \frac{1}{N} \ 1 - \frac{2}{N} \ \dots \ 0 \right]^T \quad (9)$$

The remaining eigenvalues are of absolute magnitude less than 1, implying that all except two of the diagonal elements in the eigenvalue matrix Λ^k tend to zero as $k \rightarrow \infty$. Hence, in order to obtain the limiting value of the state dynamics, it remains only to utilize the eigenvectors corresponding to the unit eigenvalues to obtain the limit in (8).

Partition the eigenvector matrices as follows

$$E = [e_1 \ e_2 \ \dots \ e_{N+1}], \quad E^{-1} = [\tilde{e}_1 \ \tilde{e}_2 \ \dots \ \tilde{e}_{N+1}]^T$$

and let e_1 and e_{N+1} denote the eigenvectors corresponding to the unit eigenvalues, i.e., $e_1 = \gamma$, $e_{N+1} = \mathbf{1} - \gamma$. It is now possible to verify from the structure of Γ and the partitioning of E and E^{-1} that

$$\tilde{e}_1^T = [1 \ 0 \ \dots \ 0], \quad \tilde{e}_{N+1}^T = [0 \ 0 \ \dots \ 1], \quad (10)$$

Finally, substituting $e_1 = \gamma$, $e_{N+1} = \mathbf{1} - \gamma$ into (8) we get

$$\lim_{k \rightarrow \infty} \Gamma^k = [\gamma \ 0 \ \dots \ 0 \ \mathbf{1} - \gamma]. \quad (11)$$

Thus it is seen that the asymptotic limit of the state is dependent only on the first and last elements of the state vector \tilde{x} . As these are fixed at 0 and T respectively, the statement of the theorem follows as desired. \triangleleft

It is of interest to note that if Γ had been a positive (or irreducible) matrix, the Perron-Frobenius theorem would have been an elegant approach to obtain the steady state matrix directly. However, in this case, owing to the zero entries and the lack of strong connectivity of the graph associated with Γ (a necessary and sufficient condition for irreducibility), the above more involved proof is required.

Finally, we state the full desynchronization algorithm as Algorithm 1, where we use the notation $rand \in [0, T]$ to denote the operation of picking a random value from a uniform distribution over the interval $[0, T]$.

4. SIMULATION EXAMPLES

Armed with Algorithm 1, we re-visit the simulation examples in Section 2. Figure 6 shows the temperature curves of a subset of a population of 10,000 identical TCLs, simulated under the same conditions as above, i.e., $R = 2^\circ\text{C/kW}$, $C = 5 \text{ kWh}/^\circ\text{C}$, $P = 14 \text{ kW}$, $\eta = 2.5$, $\theta_{sp} = 20^\circ\text{C}$, $\Delta = 1^\circ\text{C}$ and $\theta_\infty = 28^\circ\text{C}$. The top subplot shows the behavior without desynchronization, while the bottom plot shows the behavior with Algorithm 1 applied. The temperature curves clearly look more ‘jumbled’ after the step, which indicates a greater degree of desynchronization, and the operation intervals can be seen to return to their original size after in the space of a few hours.

Figure 7 shows the corresponding power curves. The effect of the desynchronization is very obvious almost immediately after the step, as the amplitude of already the first peak is less than the case without desynchronization, and most of the oscillations are suppressed after about

Algorithm 1: Desynchronize Cooling TCL

Data: Broadcast signal indicating a setpoint change δ ; Initial operation setpoint θ_{sp} and interval width Δ ; initial on/off status $s \in \{0, 1\}$; on/off period time T ; constant $a > 0$

Result: t_k yielding maximal desynchronization for $k \rightarrow \infty$

begin

```

 $\theta_{sp} \leftarrow \theta_{sp} + \delta$ ;
 $t_0 \leftarrow rand \in [0, T]$ ;
 $\alpha \leftarrow rand \in [0, \Delta/2]$ ;
 $k \leftarrow 0$ ;
while  $k < \infty$  do
  for  $0 < t < T$  do
    Measure  $\theta(t)$ ;
    if  $t < t_k$  then
       $t_{prev} \leftarrow$  time of last transition before  $t_n$ ;
    if  $t = t_k$  then
      if  $s = 0$  then
         $s \leftarrow 1$ ;
      else
         $s \leftarrow 0$ ;
    if  $t > t_k$  then
       $t_{prev} \leftarrow$  time of first transition after  $t_n$ ;
    if  $\theta(t) \leq \theta_{sp} - \frac{\Delta}{2} + \alpha$  then
       $s_i(t) \leftarrow 0$ ;
    else if  $\theta(t) \geq \theta_{sp} + \frac{\Delta}{2} - \alpha$  then
       $s_i(t) \leftarrow 1$ ;
    if  $s = 1$  then
      Turn cooling power on;
       $\alpha \leftarrow \alpha e^{-a\epsilon}$ ;
       $t \leftarrow t + \epsilon$ ;
   $t_{k+1} \leftarrow (t_{prev} + t_{next})/2$ ;
   $k \leftarrow k + 1$ ;
   $t \leftarrow 0$ ;

```




Fig. 6. Simulation with 10,000 identical TCLs subjected to a step change at time $t = 10$ h. Top: no desynchronization; bottom: with synchronization

one period. We also notice that there are some oscillations remaining, which will gradually die out due to the long-term behavior of the adaptive algorithm.




Fig. 7. Simulation with 10,000 identical TCLs subjected to a step change at time $t = 10$ h. Full line: with desynchronization; dashed: without synchronization




Fig. 8. Simulation with 10,000 inhomogeneous TCLs subjected to a step change at time $t = 10$ h. Full line: with desynchronization; dashed: without synchronization

Figure 8 shows a simulation with a heterogeneous population of TCLs. Here, as opposed to Figure 7, the parasitic oscillations die out in the case without desynchronization as well. However, Algorithm 1 clearly speeds up the process, however. Note that the oscillations that can be seen before the setpoint change is due to the initialization of the TCL on/off states, which does not match exactly with the steady-state distribution (the desynchronization is not active before $t = 10$ h). Note also that, by changing the setpoint of all the TCLs in the simulation by half a degree, it was possible to reduce the mean total power consumption by approximately 2 MW after the step.

5. DISCUSSION

This paper presented a novel algorithm for counter-acting unwanted oscillations caused by synchronization of populations of temperature controlled loads, which requires neither central management of the individual units nor communication between units. The algorithm comprises two main parts, a ‘fast’ randomization of the temperature bands within which the TCLs operate, and a ‘slow’ adaptive adjustment of enforced switch timings, which maximizes desynchronization over time. We presented a formal proof of convergence of homogeneous populations to the desynchronized status, as well as simulations that indicate that the algorithm is able to effectively dampen

power consumption oscillations for both homogeneous and heterogeneous populations.

However, the simulations also indicate that even with active desynchronization, it is hard to avoid large peaks right after a common broadcast setpoint step. Thus, high-level model-based control along the lines of the work presented in Kundu et al. [2011], will probably be preferable to simple steps.

REFERENCES

- S. Bashash and H.K. Fathy. Modeling and control of aggregate air conditioning loads for robust renewable power management. *IEEE Transactions on Control Systems Technology*, page to appear, 2012.
- D.S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(9):1389–1400, 2009.
- M. Klobasa. Analysis of demand response and wind integration in Germany’s electricity market. *IET Renewable Power Generation*, 4(1):55–63, 2010.
- S. Kundu and N. Sinitsyn. Safe protocol for controlling power consumption by a heterogeneous population of loads. In *Proc. of 2012 American Control Conference*, June 2012.
- S. Kundu, N. Sinitsyn, S. Backhaus, and I. A. Hiskens. Modeling and control of thermostatically-controlled loads. In *Proc. of the 17th Power Systems Computations Conference*, 2011.
- R.P. Malhame and C.Y. Chong. Electric load model synthesis by diffusion approximation of a higher-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control*, 30(9):854–860, 1985.
- K. Mets, T. Verschueren, W. Haerick, C. Develder, and F. De Turck. Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging. In *Network Operations and Management Symposium Workshops*, pages 293–299. IEEE, 2010.
- A. Mohsenian-Rad, v.W.S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, 2010.
- K. Moslehi and R. Kumar. A reliability perspective of the smart grid. *IEEE Transactions on Smart Grid*, 1(1): 57–64, 2010.
- C. Perfumo, E. Kofman, J.H. Braslavsky, and J.K. Ward. Load management: Model-based control of aggregate power for populations of thermostatically controlled loads. *Energy Conversion and Management*, 55(1):36–48, 2012.
- M. Petersen, J. Bendtsen, and J. Stoustrup. Optimal dispatch strategy for the agile virtual power plant imbalance compensation problem. In *Proc. of American Control Conference*, Montreal, Canada, June 2012.
- J. Short, D. G. Infield, and L. L. Freris. Stabilization of grid frequency through dynamic demand control. *IEEE Transactions on Power Systems*, 22(3):1284–1293, 2007.
- G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, 2008.